

CLAIMS

What is claimed is:

1. A method for allocating bandwidth to a switch while reducing burstiness in message traffic, the method comprising:
 - 5 (a) receiving input from a user regarding bandwidth to be allocated to at least one queue in a switch;
 - (b) converting the bandwidth into a base bandwidth value and a residual bandwidth value;
 - (c) automatically converting the base bandwidth value to a first
10 number of tokens to be placed in a token bucket associated with the queue every predetermined token bucket refresh interval;
 - (d) automatically converting the residual bandwidth value to a second number of tokens and a number c of the predetermined token bucket refresh intervals, c being an integer, wherein the
15 base bandwidth is achieved by placing the first number of tokens in the token bucket every predetermined token bucket refresh interval and wherein the residual bandwidth value is achieved by placing the second number of tokens in the token bucket every c predetermined token bucket refresh intervals; and
 - 20 (e) servicing the queue in accordance with the tokens in the token bucket.
2. The method of claim 1 wherein receiving input from a user regarding bandwidth to be allocated to at least one queue in a switch includes

receiving a bandwidth value from the user in a standard bandwidth denomination.

3. The method of claim 2 wherein the standard bandwidth denomination comprises kilobits per second.
- 5 4. The method of claim 1 wherein converting the bandwidth into a base value and a residual value includes calculating a first bandwidth value achieved by placing one token in the token bucket every predetermined token bucket refresh interval, and selecting the base and residual bandwidth values based on the first bandwidth value.
- 10 5. The method of claim 4 wherein selecting the base and residual bandwidth values based on the first bandwidth value includes setting the base bandwidth value to a portion of the bandwidth value requested by the user that is greater than the first bandwidth value and setting the residual bandwidth value to a portion of the bandwidth entered by the user that is less than the first bandwidth value.
- 15 6. The method of claim 1 wherein the predetermined token bucket refresh interval is equal to a first value being at least the number of clock cycles required to update token buckets associated with a predetermined set of the queues in the switch, thereby providing the base bandwidth value to the predetermined set of queues.
- 20 7. The method of claim 6 wherein the first value is greater than the number of clock cycles required to update the token buckets associated with the predetermined set of queues in the switch by a predetermined number of additional clock cycles, and wherein the additional clock cycles are

- used to refresh token buckets associated with a portion of the predetermined set of queues every predetermined token bucket refresh interval with the second number of tokens so that all of the token buckets associated with the predetermined set of queues will be updated with the second number of tokens every c predetermined token bucket refresh intervals, thereby providing the residual bandwidth value to each queue in the predetermined set of queues.
- 5
8. The method of claim 1 wherein the first and second numbers of tokens, the predetermined token bucket refresh interval, and c are configurable by the user.
- 10
9. The method of claim 1 wherein automatically converting the base and residual bandwidth values includes automatically converting the bandwidth values using software and using the converted values to program packet scheduling hardware.
- 15
10. The method of claim 1 wherein automatically converting the residual bandwidth to a second number of tokens and a number c of token minimum token bucket refresh intervals includes setting the second number of tokens to a predetermined fixed value and calculating the number c that achieves the residual bandwidth.
- 20
11. The method of claim 10 wherein calculating the number c includes calculating the number c using the following equation:

$$c = 8 * f / (M * I),$$

where f is the clock frequency in clock cycles per second, M is the residual bandwidth in bits per second, and I is the predetermined token bucket refresh interval in clock cycles.

12. The method of claim 1 wherein the base bandwidth value is on the order of 1 megabit per second (Mbps) and the residual bandwidth value is on the order of less than 1 Mbps.
13. The method of claim 1 wherein the base bandwidth value is on the order of 1 gigabit per second (Gbps) and the residual bandwidth value is on the order of less than 1 Gbps.
14. The method of claim 1 wherein servicing the queue in accordance with the tokens in the token bucket includes scheduling the queue based on a bandwidth state associated with the queue.
15. The method of claim 14 wherein the bandwidth state indicates whether the bandwidth allocated to the queue has been exceed.
16. The method of claim 14 wherein scheduling the queue based on the bandwidth state includes scheduling the queue based on a combination of bandwidth state, a priority assigned to the queue with respect to other queues in the switch and the relative time since the queue has been serviced with respect to other queues in the switch.
17. A method for allocating bandwidth to a queue in a switch, the method comprising:
 - (a) receiving, from a user, a desired bandwidth in a standard bandwidth denomination to be provided by a switch;

- (b) automatically converting the desired bandwidth to a token bucket refresh rate;
 - (c) refreshing at least one token bucket associated with the switch at the token bucket refresh rate; and
 - 5 (d) scheduling at least one queue in the switch to be serviced based on available tokens in the token bucket.
- 18. The method of claim 17 wherein receiving a desired bandwidth in a standard bandwidth denomination to be provided by a switch includes receiving a desired bandwidth value from the user in a domination
10 comprising bits per second and wherein automatically converting the desired bandwidth to a token bucket refresh rate includes automatically converting the bandwidth in bits per second into the token bucket refresh rate.
- 19. The method of claim 17 wherein receiving a desired bandwidth to be
15 provided by a switch includes receiving input from a user regarding minimum and maximum bandwidth values to be provided by a switch and wherein automatically converting the desired bandwidth to a token bucket refresh rate includes converting the minimum and maximum bandwidth values to first and second token bucket refresh rates at which
20 minimum and maximum token buckets associated with the switch will be refreshed.
- 20. The method of claim 17 wherein automatically converting the desired bandwidth to a token bucket refresh rate includes automatically writing the token bucket refresh rate into hardware associated with the queue.

21. The method of claim 17 wherein automatically converting the desired bandwidth to a token bucket refresh rate includes converting the desired bandwidth into a base bandwidth value and a residual bandwidth value, computing a first token bucket refresh rate corresponding to the base bandwidth value and a second token bucket refresh rate corresponding to the residual bandwidth value.
22. The method of claim 21 wherein computing the first token bucket refresh rate includes computing a first number of tokens to be placed in the token bucket every token bucket refresh interval in order to provide the base bandwidth value.
23. The method of claim 21 wherein computing the second token bucket refresh rate includes computing a number of token bucket refresh intervals over which a second number of tokens must be placed in the token bucket to provide the residual bandwidth value.
24. The method of claim 23 wherein the second number of tokens comprises a single token, thereby achieving a minimum bandwidth resolution.
25. The method of claim 23 wherein refreshing at least one token bucket includes refreshing all of the token buckets in the switch with the base token value every token bucket refresh interval and refreshing predetermined token buckets in the switch with the second number of tokens every predetermined number of token bucket refresh intervals.

26. The method of claim 17 wherein scheduling at least one queue in the switch to be serviced based on available tokens in the token bucket includes scheduling the queue based on bandwidth state of the queue.
27. The method of claim 26 wherein scheduling the queue based on bandwidth state of the queue includes assigning a bandwidth state to the queue based on whether bandwidth consumed by the queue is less than a minimum bandwidth value, between minimum and maximum bandwidth values, or greater than the maximum bandwidth value.
28. The method of claim 26 wherein scheduling the queue based on bandwidth state of the queue includes scheduling the queue based on bandwidth state, relative priority of the queue with respect to other queues in the switch, and relative time that the queue has been serviced as compared to the other queues in the switch.
29. A system for fine grain bandwidth allocation in a switched network element without burstiness, the system comprising:
- (a) a user interface operatively associated with the switched network element for receiving input from a user regarding a desired bandwidth to be provided by the switched network element;
 - (b) a bandwidth converter operatively associated with the user interface for converting the bandwidth value received from the user into at least one token bucket refresh rate;
 - (c) a traffic manager operatively associated with the bandwidth converter for maintaining a plurality of queues for scheduling packets to be forwarded from the switched network element, for

maintaining at least one token bucket for each queue, and for refreshing the token buckets at the token bucket refresh rate; and

- (d) a scheduler operatively associated with the traffic manager for scheduling the queues to be serviced based on tokens present in the token bucket for each queue.

5 30. The system of claim 29 wherein the traffic manager is adapted to maintain first and second token buckets associated with each queue in the switched network element, the first token bucket being adapted to control a minimum bandwidth to be allocated to each queue and the
10 second token bucket being adapted to control a maximum bandwidth to be allocated to each queue.

31. The method of claim 30 wherein the user interface is adapted to receive the minimum and maximum bandwidth values from the user in a standard bandwidth denomination.

15 32. The method of claim 31 wherein the standard bandwidth denomination comprises bits per second.

33. The system of claim 31 wherein the bandwidth converter is adapted to convert the bandwidth values from the standard bandwidth denomination to token bucket refresh rates.

20 34. The system of claim 29 wherein the bandwidth converter is adapted to convert bandwidth in kilobits per second into token bucket refresh rates.

35. The system of claim 29 wherein the bandwidth converter is adapted to convert the bandwidth value into base token value to be input in the token buckets every predetermined token bucket refresh interval and a

number of token bucket refresh intervals in which a residual token value is to be placed into each of the token buckets.

36. The system of claim 29 wherein the scheduler is adapted to schedule the queues based on bandwidth state of the queues.
- 5 37. The system of claim 36 wherein the scheduler is adapted to schedule the queues based on bandwidth state, relative priority of each queue in the switch, and relative time since each queue was last serviced.